

## CLAIMS

What is claimed is:

- 1 1. A method for determining the source of audio degradation in an IP telephony  
2 environment, the method comprising the computer-implemented steps of:  
3 for each switching device of a set of switching devices that are configured on a  
4 network between a first endpoint and a second endpoint,  
5 transmitting, from an originating device to the switching device, a reference  
6 version of an audio waveform;  
7 receiving a second version of the audio waveform that represents the  
8 waveform after transmission at least to the switching device;  
9 computing audio degradation that occurred between the originating device and  
10 the switching device by comparing the reference version of the audio  
11 waveform with the second version of the audio waveform; and  
12 determining the source of audio degradation in a path between the first endpoint and  
13 the second endpoint based on the audio degradations associated with each of  
14 the switching devices of the set of switching devices.
  
- 1 2. The method of claim 1, wherein the step of computing audio degradation includes  
2 computing audio degradation with a perceptual measure.
  
- 1 3. The method of claim 1, wherein the step of computing audio degradation includes  
2 computing audio degradation using Perceptual Evaluation of Speech Quality (PESQ)  
3 techniques.

1 4. The method of claim 1, wherein the step of receiving includes receiving the second  
2 version of the audio waveform at the originating device from the switching device.

## 1 5. The method of claim 4,

2 wherein the step of receiving includes receiving the second version of the audio

3 waveform over a reliable TCP connection after the switching device

timestamps packets that it received that correspond with the reference version

5 that was transmitted by the originating device and buffers the packets before

6 sending them to the originating device over the reliable TCP connection; and

7 wherein the step of computing the audio degradation consists of computing the audio

8 degradation that occurred in a path from the originating device to the

9 switching device.

## 1 6. The method of claim 1,

2 wherein the step of receiving includes receiving the second version of the audio

3 waveform at the originating device from the switching device; and

4 wherein the step of computing the audio degradation consists of computing the audio

5 degradation that occurred in a round-trip path between the originating device

6 and the switching device.

## 1 7. The method of claim 1,

2 wherein the step of receiving includes receiving the second version of the audio

3 waveform at the switching device; and

4 wherein the step of computing includes computing the audio degradation at the

5 switching device.

1 8. The method of claim 7, further comprising the computer-implemented step of:  
2 for each switching device of a set of switching devices, receiving from the switching  
3 device a measure of the audio degradation that was computed at the switching  
4 device.

1 9. The method of claim 1, further comprising the computer-implemented step of:  
2 determining the set of network switching devices that are configured on the network  
3 between the first endpoint and the second endpoint.

1 10. The method of claim 1, further comprising the computer-implemented step of:  
2 determining the set of network switching devices that are configured on the network  
3 between the first endpoint and the second endpoint by determining a set of  
4 network switching devices that are in a path from the first endpoint to the  
5 second endpoint.

1 11. The method of claim 10, further comprising the computer-implemented step of:  
2 determining the set of network switching devices that are configured on the network  
3 between the first endpoint and the second endpoint by determining the set of  
4 network switching devices that are configured in a path from the second  
5 endpoint to the first endpoint.

1 12. The method of claim 11, wherein the path from the first endpoint to the second  
2 endpoint is different than the path from the second endpoint to the first endpoint.

1 13. The method of claim 1, wherein the set of switching devices comprises every  
2 switching device configured on a path between the first endpoint and the second  
3 endpoint.

1 14. The method of claim 1, wherein the set of switching devices comprises a subset of the  
2 switching devices configured on a path between the first endpoint and the second  
3 endpoint.

1 15. The method of claim 1, wherein the step of transmitting is performed for each  
2 switching device of the set of switching devices in a sequential order from a  
3 switching device that is nearest to the first endpoint to a switching device that is  
4 nearest to the second endpoint.

1 16. The method of claim 1, wherein the steps are performed by one or more intermediate  
2 network switching devices.

1 17. The method of claim 1, wherein the steps are performed by one or more network  
2 routers.

1 18. A method for determining audio degradation in a path of an IP telephony  
2 environment, the method comprising the computer-implemented steps of:  
3 receiving from an originating device, at a network switching device, a reference  
4 version of an audio waveform; and  
5 transmitting to the originating device a second version of the audio waveform that  
6 represents the audio waveform after transmission at least to the network  
7 switching device, for computing audio degradation that occurred in a path

between the originating device and the network switching device by comparing the reference version with the second version.

1 19. The method of claim 18, wherein the steps are performed by an intermediate network  
2 switching device.

1 20. The method of claim 18, wherein the steps are performed by a network router.

1 21. The method of claim 18, wherein the reference version is received as a series of  
2 packets, the method further comprising the computer-implemented steps of:  
3 attaching a timestamp to each of the packets of the series of packets;  
4 buffering each of the packets of the series of packets;  
5 wherein the step of transmitting comprises, upon receiving the complete series of  
6 packets that represent the reference version, transmitting the second version to  
7 the originating device over a reliable TCP connection, for computing the  
8 audio degradation that occurred in a path from the originating device to the  
9 switching device.

1 22. The method of claim 21, wherein the steps are performed by an intermediate network  
2 switching device.

3 23. The method of claim 21, wherein the steps are performed by a network router.

1 24. A computer-readable medium carrying one or more sequences of instructions for  
2 determining the source of audio degradation in an IP telephony environment, which  
3 instructions, when executed by one or more processors, cause the one or more  
4 processors to carry out the steps of:

5 for each switching device of a set of switching devices that are configured on a  
6 network between a first endpoint and a second endpoint,  
7 transmitting, from an originating device to the switching device, a reference  
8 version of an audio waveform;  
9 receiving a second version of the audio waveform that represents the  
10 waveform after transmission at least to the switching device;  
11 computing audio degradation that occurred between the originating device and  
12 the switching device by comparing the reference version of the audio  
13 waveform with the second version of the audio waveform; and  
14 determining the source of audio degradation in a path between the first endpoint and  
15 the second endpoint based on the audio degradations associated with each of  
16 the switching devices of the set of switching devices.

- 1 25. The computer-readable medium of claim 24, wherein the instructions cause the one or  
2 more processors to carry out the step of computing audio degradation by computing  
3 audio degradation with a perceptual measure.
- 1 26. The computer-readable medium of claim 24, wherein the instructions cause the one or  
2 more processors to carry out the step of computing audio degradation by computing  
3 audio degradation using Perceptual Evaluation of Speech Quality (PESQ) techniques.
- 1 27. The computer-readable medium of claim 24, wherein the instructions cause the one or  
2 more processors to carry out the step of receiving by receiving the second version of  
3 the audio waveform at the originating device from the switching device.

1 28. The computer-readable medium of claim 27, wherein the instructions cause the one or  
2 more processors to carry out  
3 the step of receiving by receiving the second version of the audio waveform over a  
4 reliable TCP connection after the switching device timestamps packets that it  
5 received that correspond with the reference version that was transmitted by  
6 the originating device and buffers the packets before sending them to the  
7 originating device over the reliable TCP connection; and  
8 the step of computing the audio degradation by computing the audio degradation that  
9 occurred in a path from the originating device to the switching device.

1 29. The computer-readable medium of claim 24, wherein the instructions cause the one or  
2 more processors to carry out  
3 the step of receiving by receiving the second version of the audio waveform at the  
4 originating device from the switching device; and  
5 the step of computing the audio degradation by computing the audio degradation that  
6 occurred in a round-trip path between the originating device and the switching  
7 device.

1 30. The computer-readable medium of claim 24, wherein the instructions cause the one or  
2 more processors to carry out  
3 the step of receiving by receiving the second version of the audio waveform at the  
4 switching device; and  
5 the step of computing by computing the audio degradation at the switching device.

1 31. The computer-readable medium of claim 30, wherein the instructions cause the one or  
2 more processors to carry out:  
3 for each switching device of a set of switching devices, receiving from the switching  
4 device a measure of the audio degradation that was computed at the switching  
5 device.

1 32. The computer-readable medium of claim 24, wherein the instructions cause the one or  
2 more processors to carry out:  
3 determining the set of network switching devices that are configured on the network  
4 between the first endpoint and the second endpoint.

1 33. The computer-readable medium of claim 24, wherein the instructions cause the one or  
2 more processors to carry out:  
3 determining the set of network switching devices that are configured on the network  
4 between the first endpoint and the second endpoint by determining a set of  
5 network switching devices that are in a path from the first endpoint to the  
6 second endpoint.

1 34. The computer-readable medium of claim 33, wherein the instructions cause the one or  
2 more processors to carry out:  
3 determining the set of network switching devices that are configured on the network  
4 between the first endpoint and the second endpoint by determining the set of  
5 network switching devices that are configured in a path from the second  
6 endpoint to the first endpoint.

1 35. The computer-readable medium of claim 34, wherein the path from the first endpoint  
2 to the second endpoint is different than the path from the second endpoint to the first  
3 endpoint.

1 36. The computer-readable medium of claim 24, wherein the set of switching devices  
2 comprises every switching device configured on a path between the first endpoint and  
3 the second endpoint.

1 37. The computer-readable medium of claim 24, wherein the set of switching devices  
2 comprises a subset of the switching devices configured on a path between the first  
3 endpoint and the second endpoint.

1 38. The computer-readable medium of claim 24, wherein the instruction cause the one or  
2 more processors to carry out the step of transmitting by transmitting to each switching  
3 device of the set of switching devices in a sequential order from a switching device  
4 that is nearest to the first endpoint to a switching device is that is nearest to the  
5 second endpoint.

1 39. The computer-readable medium of claim 24, wherein the steps are performed by one  
2 or more intermediate network switching devices.

1 40. The computer-readable medium of claim 24, wherein the steps are performed by one  
2 or more network routers.

1 41. A computer-readable medium carrying one or more sequences of instructions for  
2 determining audio degradation in a path of an IP telephony environment, which

3       instructions, when executed by one or more processors, cause the one or more  
4       processors to carry out the steps of:  
5       receiving from an originating device, at a network switching device, a reference  
6               version of an audio waveform; and  
7       transmitting to the originating device a second version of the audio waveform that  
8               represents the audio waveform after transmission at least to the network  
9               switching device, for computing audio degradation that occurred in a path  
10          between the originating device and the network switching device by  
11               comparing the reference version with the second version.

1   42. The computer-readable medium of claim 41, wherein the steps are performed by one  
2       or more processors configured on an intermediate network switching device.

1   43. The computer-readable medium of claim 41, wherein the steps are performed by one  
2       or more processors configured on a network router.

1   44. The computer-readable medium of claim 41, wherein the reference file is received as  
2       a series of packets, and wherein the instructions cause the one or more processors to  
3       carry out the steps of:  
4       attaching a timestamp to each of the packets of the series of packets;  
5       buffering each of the packets of the series of packets;  
6       wherein the step of transmitting comprises, upon receiving the complete series of  
7       packets that represent the reference version, transmitting the second version to  
8       the originating device over a reliable TCP connection, for computing the

9 audio degradation that occurred in a path from the originating device to the  
10 switching device.

1 45. The computer-readable medium of claim 41, wherein the steps are performed by one  
2 or more processors configured on an intermediate network switching device.

1 46. The computer-readable medium of claim 41, wherein the steps are performed by one  
2 or more processors configured on a network router.

1 47. A system for determining the source of audio degradation in an IP telephony  
2 environment, the system comprising:

3 means for transmitting a reference version of an audio waveform from an originating  
4 device to each switching device of a set of switching devices that are  
5 configured on a network between a first endpoint and a second endpoint;  
6 means for receiving a second version of the audio waveform that represents the  
7 waveform after transmission at least to each respective switching device of the

8 set of switching devices;  
9 means for computing audio degradation that occurred between the originating device  
10 and each switching device of the set of switching devices by comparing the  
11 reference version of the audio waveform with the second version of the audio  
12 waveform; and

means for determining the source of audio degradation in a path between the first endpoint and the second endpoint based on the audio degradations associated with each of the switching devices of the set of switching devices.

1 48. The system of claim 47, wherein the system is configured on a network router.

1 49. A system for determining audio degradation in a path of an IP telephony  
2 environment, the system comprising:  
3 means for receiving from an originating device, at a network switching device, a  
4 reference version of an audio waveform; and  
5 means for transmitting to the originating device a second version of the audio  
6 waveform that represents the audio waveform after transmission at least to the  
7 network switching device, for computing audio degradation that occurred in a  
8 path between the originating device and the network switching device by  
9 comparing the reference version with the second version.

1 50. The system of claim 49, wherein the system is configured on a network router.

1 51. A network device that can determine the source of audio degradation in an IP  
2 telephony environment, the device comprising:  
3 a network interface;  
4 a processor coupled to the network interface and receiving messages from a network  
5 through the network interface;  
6 a computer-readable medium comprising one or more stored sequences of  
7 instructions which, when executed by the processor, cause the processor to  
8 carry out the steps of:  
9 for each switching device of a set of switching devices that are configured on  
10 a network between a first endpoint and a second endpoint,  
11 transmitting, from an originating device to the switching device, a  
12 reference version of an audio waveform;

13 receiving a second version of the audio waveform that represents the  
14 waveform after transmission at least to the switching device;  
15 computing audio degradation that occurred between the originating  
16 device and the switching device by comparing the reference  
17 version of the audio waveform with the second version of the  
18 audio waveform; and  
19 determining the source of audio degradation in a path between the first  
20 endpoint and the second endpoint based on the audio degradations  
21 associated with each of the switching devices of the set of switching  
22 devices.

1 52. A network device that can determine audio degradation in a path of an IP telephony  
2 environment, the device comprising:  
3 a network interface;  
4 a processor coupled to the network interface and receiving messages from a network  
5 through the network interface;  
6 a computer-readable medium comprising one or more stored sequences of  
7 instructions which, when executed by the processor, cause the processor to  
8 carry out the steps of:  
9 receiving from an originating device, at a network switching device, a  
10 reference version of an audio waveform; and  
11 transmitting to the originating device a second version of the audio waveform  
12 that represents the audio waveform after transmission at least to the  
13 network switching device, for computing audio degradation that

14                   occurred in a path between the originating device and the network  
15                   switching device by comparing the reference version with the second  
16                   version.